

Political Disaffection: a case study on the Italian Twitter community

C. Monti¹, A. Rozza², G. Zappella³, M. Zignani¹, A. Ardvissan⁴, M. Poletti⁴

¹ Dipartimento di Informatica, Università degli Studi di Milano
{monti, zignani}@di.unimi.it

² Dipartimento di Scienze Applicate, Università degli Studi di Napoli “Parthenope”
alessandro.rozza@uniparthenope.it

³ Dipartimento di Matematica, Università degli Studi di Milano
giovanni.zappella@unimi.it

⁴ Dipartimento di Scienze Politiche, Università degli Studi di Milano
{adam.ardvissan, monica.poletti}@unimi.it

Abstract

In our work we analyse the political disaffection or “the subjective feeling of powerlessness, cynicism, and lack of confidence in the political process, politicians, and democratic institutions, but with no questioning of the political regime” by exploiting Twitter data through machine learning techniques.

In order to validate the quality of the time-series generated by the Twitter data, we highlight the relations of these data with political disaffection as measured by means of public opinion surveys. Moreover, we show that important political news of Italian newspapers are often correlated with the highest peaks of the produced time-series.

1 Introduction

Twitter is one of the biggest micro-blogging services in the world. Micro-blogging refers to the publication of short text messages, used to share all kinds of information; on Twitter, these messages are called “tweets” (their maximum length is 140 characters), and many millions of them are posted every day.

Twitter is an interesting data source to explore public sentiment trends (Bollen, Mao, and Pepe (2011); Vallina-Rodriguez et al. (2012)): its content is easily available, and it has a very flexible nature due to the fact that it is currently used for open conversations, public opinions, and news commentaries. Another crucial characteristic of Twitter content is its timeliness; this peculiarity guarantees that tweets are related to a much closer temporal window with respect to other user-generated texts, such as blogs.

Modelling trends from Twitter data has become a popular research task. Among such studies, those drawing attention to political topics are some of the most attractive, and in the last years a great deal of research has focused on them.

In this study we want to concentrate on an important concept in political science: *political disaffection*, or “the subjective feeling of powerlessness, cynicism, and lack of confidence in the political process, politicians, and democratic institutions, but with no questioning of the political

regime” (Di Palma, 1970). In political science, levels of political disaffection are understood to relate to levels of political participation and, consequently to have important implications for the legitimacy of democratic political systems. This makes the study of political disaffection one of the key topics of contemporary studies of political behavior. Despite the popularity of the term in the relevant literature, political disaffection is a complex and multi-dimensional term. The concept of political disaffection does not necessarily imply low levels of satisfaction with the current government, nor adversity to the system of democratic politics. Instead, concept of political disaffection has two components. One, low trust in politics or politicians in general. Two, what is known as political inefficacy, “a belief that politics is complex, difficult to understand, self-referential and distant from citizens” (Campbell, Gurin, and Miller, 1954).

To our knowledge political disaffection has never been studied using Twitter data. In this work we propose an automatic approach to measuring political disaffection using Twitter data with the aim to study the relations between our measurement of political disaffection and political disaffection as measured by public opinion surveys. In accordance with the literature we define political disaffection as a general discontent with the political system, and not as a partisan position against a particular party, individual or policy, we operationalize this definition in three steps. One, we use a supervised methodology to extract a subset of political tweets from the universe of tweets in Italian. Two, we perform a sentiment analysis to analyze political tweets with negative sentiment. Three, we automatically select the tweets that refer to politics of politicians in general, rather than specific political events or personalities. In order to validate our operational measurement, the selected tweets that represent political disaffection are used to create time series that are related to indicators of political disaffection in public opinion surveys. Furthermore, we also show that important political news from Italian newspapers are often correlated with peaks in the produced time-series.

This paper is organized as follows: in Section 2 the related works are summarized; in Section 3 we describe the procedures used to generate the datasets employed to train our supervised methods, the approach to extract the tweets employed in our analysis, and we summarize the public opinion surveys used to validate the quality of our approach; in Sec-

tion 4 we describe the employed feature extraction methodologies and the classification approach that we used; in Section 5 we present the achieved results; Section 6 highlights the conclusions of our work.

2 Related Works

In literature a great deal of research has focused on the analysis of different phenomena using the data of micro-blogging services. Among them, we recall the work proposed in (Popescu and Pennacchiotti, 2011) where the authors explore the correlation between types of user engagement and events about celebrities using Twitter data. Furthermore, (Bollen and Mao, 2011) propose an approach able to predict the stock market by employing micro-blogging data.

The most closely related works are those concerning the analysis of political topics by employing Twitter data. In (Bollen, Mao, and Pepe, 2011), the authors propose a method able to extract from tweets data different time series corresponding to the evolution of 6 emotional attributes (tension, depression, anger, vigour, fatigue, and confusion) called Profile of Mood States (POMS). The authors applied POMS to suggest that socio-economic agitations caused significant fluctuations of the mood levels.

One of the earliest papers discussing the feasibility of using Twitter data as a substitute for traditional public opinion surveys has been proposed in (O'Connor et al., 2010). The authors employ Opinion-Finder¹ to determine both a positive and a negative score for each tweet in their dataset. Then, raw numbers of positive and negative tweets regarding a given topic are used to compute a sentiment score. Using this method, sentiment time series are prepared for a number of topics such as: presidential approval, consumer confidence, and US 2008 Presidential elections. According to the authors both consumer confidence and presidential approval public opinion surveys show correlation with the Twitter sentiment data computed with their approach. However, no correlation are been found between electoral public opinion surveys and Twitter sentiment data.

In (Tumasjan et al., 2010) an analysis of the tweets related to different parties running for the German 2009 Federal election is presented. Moreover the authors show that the count of tweets mentioning a party or a candidate accurately reflected the election results suggesting a possible approach to perform an electoral prediction. Furthermore, in (Livne et al., 2011) a novel method that tries to predict elections has been presented. This approach relies both on Twitter data and also in additional information such as the party a candidate belongs to, or incumbency. In (Bermingham and Smeaton, 2011) is proposed a methodology that extends the previous approaches to incorporate sentiment analysis to perform a prediction on the political election. The authors test their method in the 2011 Irish General Election finding that it is not competitive when compared with traditional public opinion surveys. Similar approaches are

¹Opinion-Finder is a system that performs subjectivity analysis, automatically identifying when opinions, sentiments, speculations and other private states are present in text.

proposed in (Tjong Kim Sang and Bos, 2012; Skoric et al., 2012).

Nevertheless, some works present some doubts about the possibility to perform an electoral prediction using Twitter data (Metaxas, Mustafaraj, and Gayo-Avello, 2011; Gayo-Avello, 2012). More precisely, in (Metaxas, Mustafaraj, and Gayo-Avello, 2011) the authors analyze the results from a number of different elections and they conclude that Twitter data is only slightly better than chance when predicting elections.

In our work we analyse the well-known political attitude of political disaffection by employing Twitter data through machine learning techniques. In order to validate the quality of the information extracted from the Twitter data, we highlight the relations of these data with political disaffection as measured in public opinion surveys. Here, the attitude of political inefficacy is used as proxy of this broad concept. This attitude expresses the (subjective) sense of powerlessness of citizens in politics. Its symmetrical concept, political efficacy, is instead the individual self-image as an influential participant in politics, “the feeling that individual political action does have, or can have, an impact upon the political process...the feeling that political and social change is possible, and that the individual citizen can play a part in bringing about this change” (Campbell, Gurin, and Miller, 1954). Political efficacy is considered crucial for participation since efficacious citizens have a greater propensity to engage in political action, and high levels of political efficacy within a population are important in creating support for the political system. If increasingly more citizens believe that they have not enough abilities to influence political decision-making and opportunities to participate in politics, whilst simultaneously not believing in the accountability of how the political system works, they will become frustrated and discontent, reducing their likelihood to act politically or to vote. Together with political inefficacy we also use low intentions to vote as a proxy of political disaffection, since we believe it is a behavioural consequence of citizens’ sense of powerlessness (see Torcal and Montero (2006)).

3 Dataset

In this section we describe how we have generated the dataset employed to train our supervised methods. Furthermore, the procedure used to extract the tweets employed to identify political disaffection is described. Finally, the information extract from different Italian newspapers, and the employed public opinion surveys are described.

Training Data

In order to train the classifiers that compose our methodology to extract the political disaffection, we build the training set employing the Twitter API v1 by a 2-step procedure involving a semi-automatic search method and a labelling phase guided by experts. The collection phase began at the beginning of April 2012 and ended at the beginning of June 2012 and collected about 120,000 of tweets and retweets. The collected tweets resulted from a geo-localized trending

topic² search and a targeted search on political themes. In particular, at the end of each day we requested the top 10 trending topics for Italy. As most trending topics regard non-political arguments (i.e. celebrities, sports or viral hashtags) we selected the political ones and a subset of the non-political. Furthermore, in order to have a more meaningful number of political tweets, we searched for tweets related to politicians, political news from online newspapers and talk-shows. As query keywords we choose the name of the most famous Italian politicians, the topics of the top news in the political section of online newspapers and the official hashtags of TV-talks. After the deletion of retweets, the resulting dataset is made up by a large corpus of 40,000 records consisting of the tweet text, its date and the keyword used in the search.

Once the dataset was collected, we started the labelling phase employing the expertise of a pool of 40 sociology and political science students. Each student was assigned a set of 3000 tweets to be classified by means of a web application. Two different labels have been associated to each tweet, the first is political/not political (the students need to identify if the tweet topic is political), and the second is positive or neutral sentiment versus negative sentiment of the political tweets. As the meaning of the labels are quite fuzzy, before the labelling process we made a kick-off meeting to ensure a label definition agreement and consequently an acquisition of homogeneous and reliable data.

The tweets' assignment has been made so that each tweet text has been labelled by three different students. This way we can increase the accuracy and the meaningfulness of the labelling process by selecting tweets for which all the evaluators agree on its political nature, and we employ a majority voting approach for labelling the sentiment of the tweets. Precisely, the sentiment label is set in according to a decision rule that selects the label which has a majority, that is, more than half the votes. Note that, taking into account the Krippendorff's alpha coefficient³, we obtain for the sentiment label $\alpha = 0.79$. The final dataset (TData) is then composed by 31,000 labelled tweets. To the best of our knowledge, it represents one of the biggest dataset containing tweets classified by experts.

Newspaper Data

The adoption of TData in the training could presents some drawbacks due to the limited period it spans. For example, some important features for a classifier could be characteristic of the retrieval period and could lost their relevance accounting for a wider period. These drawbacks result in a limited ability of generalization of the model employed to classify the political tweets. To improve its generality, we built up an additional dataset (NData) containing all the article titles of different Italian newspapers (*Repubblica*, *il Manifesto*, and *Libero*) so that they spanned the whole spec-

trum of the political points of view from the Right wing to Left one. More precisely, we selected, from the feed RSS history, all the articles from January the 1st 2012 to October the 10th 2012 extracting the news title, and we employed the categorization proposed by the newspaper to associate a label to the title. Precisely, if a news belongs to the political category proposed by the newspaper we set the label to 1, otherwise -1. The resulting NData is composed by 17,388 labelled newspaper titles, 10,670 of which political (61%).

Data employed for our Analysis

TData and NData were preliminary to train a supervised methodology which can detect the tweets useful for our goal (to identify political disaffection). To obtain general results on the political disaffection we perform our analysis on a large sample of Italian Twitter community. To achieve this goal, we randomly extract 50,000 Italian users, which have posted at least one Italian tweet⁴ in a fixed temporal range (October 10th to October 30th). Moreover, to extend this set to include also the less active users we perform a one-level snowball sampling (Atkinson and Flint, 2001); precisely, we select for each user all its Italian followers, thus producing a set of 261,313 users. Note that we have not expanded recursively this process to reduce the intrinsic bias produced by the snowball sampling process. Moreover, we take into account only the user profiles that has been created before April the 4th (obtaining 167,557 users), thus to prevent the problem of the continuous grow of the Italian Twitter community, that could affect the quality of our political disaffection investigation. Finally, from each selected user we extract, for the period of interest (April the 4th 2012 to October the 10th 2012), all its tweets, thus producing our final set composed by 35,882,423 tweets (TCorpus).

Public Opinion Surveys

The public opinion surveys have been collected by IPSOS from April the 11th 2012 to October the 10th 2012. The sampling procedure consisted in a survey through CATI (computer-assisted telephone interview) of a representative sample of the Italian electorate. More precisely, almost every week, respondents are contacted with a quota sampling on fixed parameters (age, gender, education) using the technique of random digit dialling.

We build two indicators for political disaffection. The first one (NO_VOTE) is broader and measures the intentions not to vote at the next elections, a behavioural consequence of disaffection. It includes the percentage of survey respondents that declare to have very low intention to vote at the next elections⁵ (see Table 1). In details, we consider the people that answered 1 (1 low - 10 high) to the question "How likely is it that you will vote at the next election?".

The second indicator (INEFFICACY) is more specific. It measures the attitude of political inefficacy, that is to say the disbelief in the accountability of the political system and of

²Trending topics are the most popular and talked-about words and phrases on Twitter for a specif time period.

³Krippendorff's alpha coefficient (Krippendorff, 2003) is a statistical measure of the agreement achieved when coding a set of units of analysis in terms of the values of a variable.

⁴To identify if a tweet is written in Italian we employ the Guess-Language library (<https://code.google.com/p/guess-language/>).

⁵The total sample consists in 24,971 respondents (~ 1040 respondents per poll).

t_i	INEFFICACY	NO_VOTE
2012-04-11	14.86%	13.23%
2012-04-18	13.77%	14.53%
2012-05-02	22.20%	19.05%
2012-05-09	-	13.37%
2012-05-16	12.93%	13.34%
2012-05-23	16.31%	12.47%
2012-06-05	12.07%	13.31%
2012-06-06	11.03%	13.76%
2012-06-13	-	10.99%
2012-06-20	10.77%	13.08%
2012-06-26	6.91%	9.29%
2012-06-27	6.84%	13.09%
2012-07-04	-	11.88%
2012-07-11	7.87%	10.04%
2012-07-17	9.51%	13.64%
2012-07-18	6.00%	10.03%
2012-07-25	-	13.26%
2012-09-04	-	13.53%
2012-09-12	8.46%	11.22%
2012-09-19	-	12.75%
2012-09-25	9.44%	12.04%
2012-09-26	10.46%	12.74%
2012-10-03	-	12.87%
2012-10-10	11.76%	14.38%

Table 1: Public Opinion Surveys for INEFFICACY and NO_VOTE indicators.

all political parties (see Table 1). We operationalize it using the propensity to vote (PTV) of respondents for a specific party⁶ (1 low - 10 high). Therefore, we include in this indicator the respondents that have low (equal to 1) PTV for all political parties included in the survey, or, alternatively, for those who have low (equal to 1) PTVs for some political parties and missing answers for the other ones. The PTV has been coded for all major Italian parties⁷.

4 Classification

Identifying political disaffection is a complex task even for human being, so, in order to create a system for the detection of this attitude in tweets, we have to define it in a formal way. The goal is to measure disaffection in conceptually similar ways to what is measured by public opinion surveys. To that end, in order to be labelled as an expression of political disaffection a tweet has to match the following three criteria:

- **Political:** the tweet should regard politics.
- **Negative:** the sentiment of the tweet should be negative.
- **Generic:** the message have to regard politicians or parties in general. Tweets regarding only a political party or specific politician are not considered.

⁶The total sample consists in 38,537 respondents (~ 2267 respondents per poll).

⁷ PD (Partito Democratico), PDL (Popolo della Libertà.), Lega Nord, IdV (Italia dei Valori), UDC (Unione di Centro), FLI (Futuro e Libertà.), SeL (Sinistra Ecologia Libertà.), and M5S (Movimento 5 Stelle)

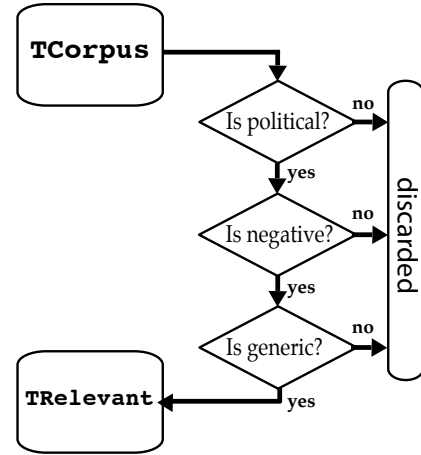


Figure 1: Classification chain employed.

Since we can not train a classifier able to consider all this criteria at the same time, we created a “chain” of classifiers as described in Figure 1. There are three fundamental steps, where the relevant tweets after each step became the input for the next one. After every step the number of relevant tweets is less or equal to the number of relevant tweets after the previous step. The relevant tweets after the third step are definitively classified as relevant and all the other tweets are classified not relevant. Roughly speaking, we give TCorpus as input of the chain and obtain a set of tweets denoting political disaffection (TRelevant) as output.

For the first step, we trained a classification algorithm using TData and NData; the resulting classifier is able to distinguish between “political” and “non-political” tweets. In the second step, the algorithm is trained with TData and the resulting classifier distinguish between tweets with negative sentiment and non-negative one. The third and last step is performed by an ad-hoc classifier created with a rule-based approach to identify the generic speech. Please note that the TData collection is fixed, but the features are extracted in different ways depending on classification step taken into account. In the next sections we summarize the feature extraction methodologies, and, subsequently, the tested classification approaches.

Feature extraction

The efficacy of textual classification crucially depends on how the textual data are transformed into numerical features. Nevertheless, identifying the best function for feature extraction and the best tokenization method is an hard problem, and the results are usually task dependent. For these reasons, we separately manage the two supervised classification tasks: politic topics and negative sentiment. Note that, in political topics we employ both the tweets data (TData) and newspaper titles (NData).

We compared different techniques for features extraction in order to find the most suitable for our problems: 5-grams

of characters, space-separated words⁸, $\{1, 2, 3\}$ -grams of words, and string kernels (Lodhi et al., 2002). Concerning the counting function, we compute: term frequency (Manning, Raghavan, and Schütze, 2008), boolean term presence (Yelena Mejova, 2011), and term frequency-inverse document frequency ($tf-idf$, Manning, Raghavan, and Schütze (2008)).

The test was performed for the “Politics” topic classification using a 4-fold cross-validation with an online linear classifier. Our results show that 5-grams of characters, independently of the counting scheme, constitutes our best option.

On the other hand, taking into account the negative sentiment problem and replicating the experiments with the same methodology of the previous case, we note that the feature characterized by the space-separated word tokenization achieves the overall best results, employing as counting scheme the term frequency. Moreover, an important improvement is given by performing a stemming process and collapsing synonyms into a single feature; to perform this task we employ a freely available Italian synonyms dictionary⁹.

An important step to ensure a robust classification is to remove the sentiment target in each tweet. To perform this task, we employed DBpedia. This database, extracted from Wikipedia, allows to perform queries and provides a simple way to capture the semantic behind words. Then, we collect a list of possible targets for the sentiment task, by obtaining via queries a full and up-to-date set of Italian political parties, politicians, and political offices. Combining this data with the recognition of strings starting with “@” (Twitter user-names) we are able to remove the sentiment target from the tweets.

Classification algorithms

We need classifiers able to scale on huge corpus possibly able to be updated over time. So, we especially focused our attention on online classifiers since they require only one sweep on the data, making the classification process really fast with really good performances on the accuracy side.

We ran all the experiments on an ordinary workstation: Intel(R) Core(TM) i7-2600K CPU at 3.40GHz with 16Gb of RAM.

We tested four different online algorithms for classification¹⁰ and one batch classification algorithm:

- **ALMA** (Gentile, 2001): is a fast classifier which try to approximate the maximal margin hyperplane between the two classes. We set the parameter p equals to 2 and we also tested different values of α .
- **OIPCAC** (Rozza et al., 2012): is a classification method employing a modified approach to estimate the Fisher Subspace, which allows to manage classification tasks

⁸Considering the space-separated words’ approach, we recognize as single word also emoticons, single punctuation marks such as “!”. The URLs are also transformed in an unique token: $\langle link \rangle$.

⁹<http://webs.racocatala.cat/llengua/it/sinonimi.htm>

¹⁰Most of these algorithms are well known and have a MATLAB implementation available in DOGMA Orabona (2009).

Classifier	Accuracy	F-Measure	Global time
ALMA	0.883 ± 0.014	0.886 ± 0.011	13.5 ± 1
PA	0.889 ± 0.012	0.890 ± 0.012	10.62 ± 0.1
PEGASOS	0.882 ± 0.010	0.883 ± 0.010	1103 ± 10
OIPCAC	0.889 ± 0.001	0.891 ± 0.010	5911 ± 52

Table 4: 10-fold results for political topic detection (in bold face the best results considering F-measure). In italic the classifier selected for the classification process. With “time”, we mean time employed for training and classification, in seconds. We were not able to conclude all the runs with RF due to its time cost.

Classifier	Accuracy	F-Measure	Global time
ALMA	0.703 ± 0.029	0.745 ± 0.034	0.82 ± 0.28
PA	0.665 ± 0.064	0.705 ± 0.124	0.91 ± 10^{-3}
PEGASOS	0.691 ± 0.033	0.732 ± 0.045	76 ± 0.1
OIPCAC	0.714 ± 0.026	0.751 ± 0.024	121 ± 25
RF	0.724 ± 0.026	0.776 ± 0.027	2173 ± 48

Table 5: 10-fold results for political topic detection (in bold face the best results considering F-measure). In italic the classifier selected for the classification process. With “time”, we mean time employed for training and classification, in seconds.

where the space dimensionality is bigger than, or comparable to, the cardinality of the training set, and to deal with unbalanced classes.

- **PASSIVE AGGRESSIVE** (PA, Crammer et al. (2006)): is a Perceptron-like method. In our experiments we test only the binary classifier with different settings.
- **PEGASOS** (Shalev-Shwartz et al., 2011): is a well-know online Support Vector Machine (SVM) solver.
- **RANDOM FOREST** (RF, Breiman (2001)): is an algorithm based on an ensemble of classification trees. Since the algorithm is widely used in machine learning challenges with good results, we will use it as yardstick in our comparison.

We compared these algorithms on the two aforementioned tasks:

- **Political**: a binary classification of tweets into “related with politics” and “not related with politics”
- **Negative**: a binary classification of tweets into “tweets with negative sentiment” and “without negative sentiment” (that is, objective or with positive sentiment).

Note that we test the online learning algorithms in a batch setting. For this particular case, we use only their one-sweep behaviour in order to speed up the classification process. After an extensive tuning of the parameters, in Table 4 and in Table 5 we report for each predictor its best performances in 10-fold cross validation on the Political classification task and on the Negative sentiment identification respectively.

In order to obtain a good trade off between accuracy and running times, we adopted a combination of PA and ALMA.

Interval	ρ	95% C.I.	P-Value for $\rho > 0$
Δ_1^{14}	0.7860	0.476-0.922	0.031%
Δ_7^{14}	0.7749	0.454-0.917	0.042%
Δ_1^7	0.6880	0.310-0.878	0.226%

Table 2: Pearson correlation index achieved between Twitter political disaffection and `INEFFICACY` time-series.

Interval	ρ	95% C.I.	P-Value for $\rho > 0$
Δ_1^{14}	0.5920	0.248-0.803	0.231%
Δ_7^{14}	0.5579	0.190-0.788	0.567%
Δ_1^7	0.4433	0.049-0.718	3.00%

Table 3: Pearson correlation index achieved between Twitter political disaffection and `NO_VOTE` time-series.

In particular we used `PA` for the political/non-political classification and `ALMA` for the sentiment classification part. It is also possible to obtain comparable performances with different combinations of classifiers.

It is important to underline that since the “Generic” task (see Section 4) is trivial, we opt to employ an approach based on a list of a few words selected by domain experts to solve it.

5 Results

In this section we describe the time-series obtained employing the informations extracted with the approach describes in Section 4 and the relations between them and the public opinion surveys. Moreover, we summarize our methodology to identify the political news that produce the highest peaks of the generated time-series (breaking news).

To perform a correlation analysis with the `INEFFICACY` indicator taken from surveys, we employ the approach described in Section 4 to generate the set of tweets denoting political disaffection (`TRelevant`). Subsequently, taking into account each survey sampling date t_i (see Table 1), we generate three time-series computing the ratio between the number of political disaffection tweets and the number of political ones by employing three time intervals:

1. from the date of the survey to 14 days before (Δ_1^{14});
2. from the day of the survey to 7 days before (Δ_1^7);
3. from 7 days before the date of the survey to 14 before (Δ_7^{14}).

Note that the same approach has been employed for the `NO_VOTE` indicator (also taken from the suveys).

In Table 2 it is shown the Pearson correlation index computed between the political disaffection tweet-series and the `INEFFICACY` time-series. It is possible to underline that, the best result (0.79) represents a strong correlation value between `INEFFICACY` and the information extracted by our methodology. Furthermore, it is important to highlight that the best time interval is Δ_1^{14} . This result suggests that Twitter is able to capture the change in citizens political disaffection more promptly than what public opinion surveys are able to do (as can be seen in Figure 2).

In Table 3 it is shown the Pearson correlation index computed between the political disaffection tweet-series and the `NO_VOTE` one. As can be noticed, this results (0.59) represents a medium correlation value but it is important because it shows that there is some connection between the modelled political disaffection and the intention to not participate at the next election day.

Taking into consideration specific affordances of twitter as a medium (i.e. immediate diffusion of news and information), these results indicate that the data that we have extracted from Twitter and data derived from surveys are two reflections of a common underlying development that exhibit different temporal characteristics. This leads us to suggest that Twitter data could be taken as a valid alternative measurement of political disaffection.

Breaking News Identification

Having verified the correlation between `INEFFICACY` and the Twitter political disaffection, we have employed the `TRelevant` data to empirically determine some of the possible causes that produce the variation in disbelief in politics and politicians, hypothesising that citizens political inefficacy is affected by controversial political news reported daily in media. To achieve this goal we have identified the peaks of the time-series generated as the daily ratio between the number of political disaffection tweets and the number of political ones, and we have associated to each peak a news belonging to `NData`.

More precisely, to identify the peaks we have employed an approach similar to that proposed in (Gruhl et al., 2004), taking into account as peak the points of the time series greater than $\mu + 2\sigma$, where μ is the mean of the points of the time-series and σ is the standard deviation; however to improve the quality of our results we have considered for each point a set of its neighbourhood¹¹ (instead of all the points) to estimate the local μ s and σ s. The qualitative results are shown in Figure 3.

To associate to each peak a news, firstly we have created an inverse document frequency (`idf`) map by employing the words extracted from the corpus of the news included in

¹¹We use a temporal window of 10 days, 5 before and 5 after the point of the time-series taken into account.

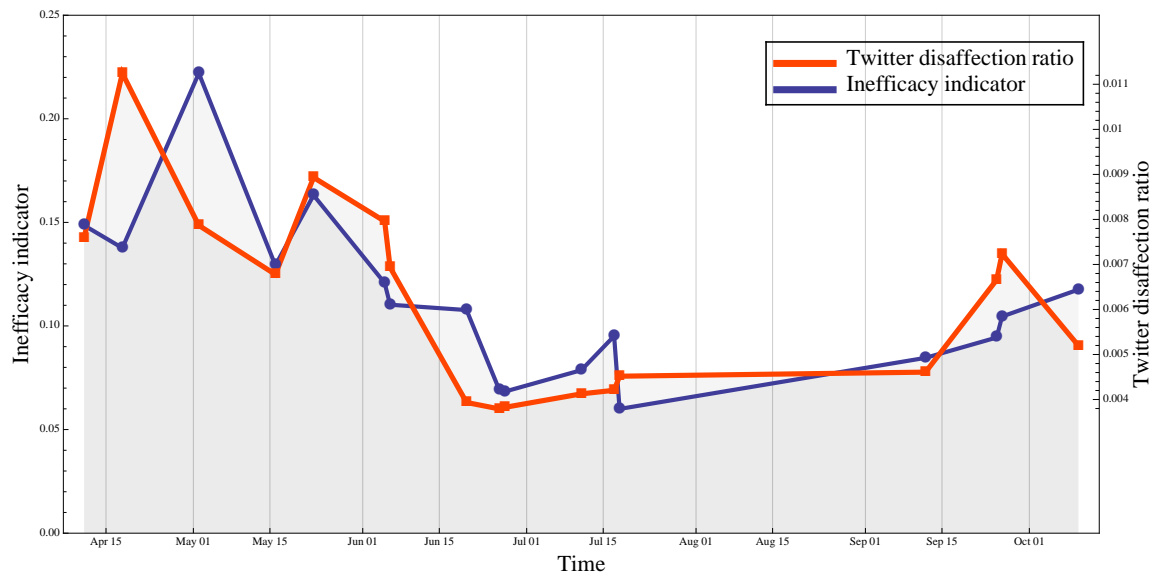


Figure 2: Twitter political disaffection time-series employing Δ_1^7 compared with the INEFFICACY indicator.

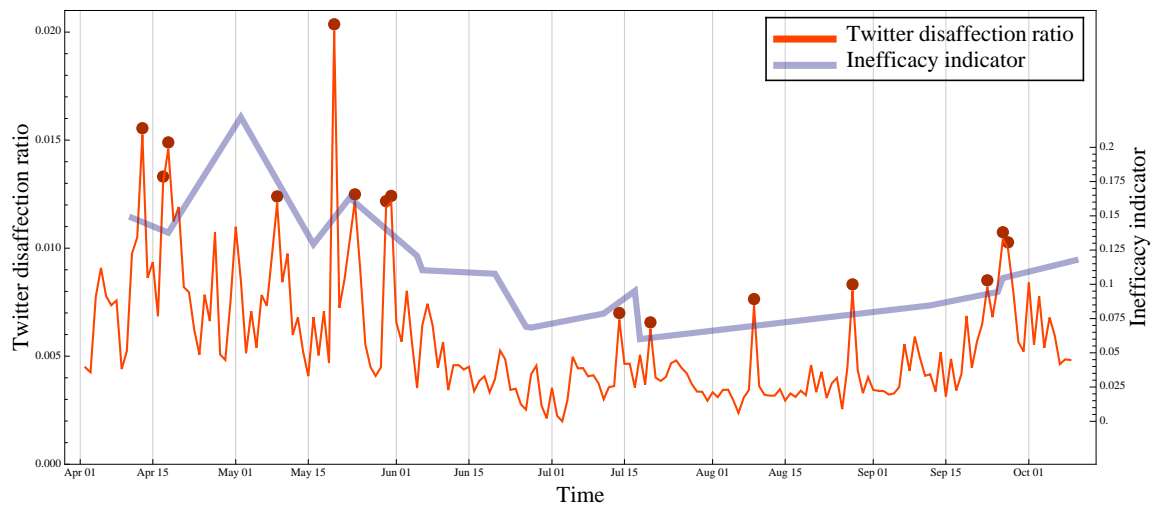


Figure 3: Political disaffection tweets day by day, with the selected peaks (highlighted by circles).

2012-04-13	✓ La Lega prova a rifarsi un'immagine. Rinuncia agli ultimi rimborsi elettorali. Lega tries to clear its name. It opts last electoral refunds out.
2012-04-17	✓ Lavoro, Monti pensa alla fiducia "I partiti approveranno la riforma". Labor, Monti thinks of a vote of confidence "Parties will enact reform".
2012-04-18	✓ Monti: niente crescita fino al 2013, disagio lavoro per meta famiglie. Monti: no economic growth until 2013, disadvantage for half of families.
2012-05-09	✓ Bersani: "Pd più forte, Monti ci ascolti" Grillo: "Partiti morti". Crollo del Pdl. Bersani: "PD stronger, Monti listen to us" Grillo: "Parties are dead". PDL falls.
2012-05-20	✓ Grillo su Brindisi: strage di Stato, fa comodo a loro. Grillo about Brindisi: state massacre, it's convenient for them.
2012-05-24	× Grillo attacca: "Noi soldi non li vogliamo rinunceremo a rimborsi prossime politiche". Grillo bashes: "We don't need money, we opt last electoral refunds out".
2012-05-30	× Riforma Csm, il gelo di Monti così é fallito il piano di Ctrialá. CSM reform, Monti's chill, Ctrialá's project is doomed.
2012-05-31	× Spread, Monti resta preoccupato "Rischio contagio malgrado gli sforzi". Spread, Monti worried "Risk contagion despite the efforts".
2012-07-14	✓ Cicchitto: "Primarie sono inutili Berlusconi candidato premier". Cicchitto: "Primary election is useless, Berlusconi is the premier candidate".
2012-07-19	× Monti ora teme il crac della Sicilia. Now Monti is afraid of Sicily default.
2012-08-09	✓ Monti al Wsj: "Con Berlusconi spread a 1200" L'ira del Pdl. E votano contro il governo. Monti to WSJ: "Spread at 1200 with Berlusconi" PDL anger and vote against government.
2012-08-28	✓ Grillo a Bersani: "Io fascista? Tu sei un fallito d'accordo con la P2". Grillo to Bersani: "Am I Fascist? You're failed at one with P2".
2012-09-23	✓ Così si rubava alla Regione Lazio ecco le rivelazioni di Fiorito ai pm. Fiorito's admission to public prosecutor, how they stole at Lazio government.
2012-09-26	✓ Caso Sallusti, salta l'accordo con il giudice il direttore domani rischia il carcere. Sallusti's instance, legal agreement breaks, the lead director risks the jail.
2012-09-27	✓ Sallusti: "In Italia mancano le palle" Paolo Berlusconi respinge le dimissioni. Sallusti: "In Italy many wimps" Paolo Berlusconi rejects Sallusti's resignation.

Table 6: Each row represents the news with the highest cosine similarity identified by our approach. The symbol ✓ represents the identified news that also appears in the political Twitter trending topic of the day taken into account. As additional contextual information, Lega is an Italian party, Bersani, Cicchitto and Berlusconi are politicians, PD is the Italian Democratic Party, Monti is the Italy premier, Sallusti and Paolo Berlusconi (Silvio Berlusconi's brother) are respectively the lead director and the editor of a newspaper, Grillo is a comic/politician, Fiorito is a regional councilman involved in bribe inquiry. CSM is the magistrates' internal board of supervisors.

NData, and 1,000,000 of tweets (PTCorpus) randomly selected from the political subset of TCorpus (we employ the same classifier used for the political task as described in Section 4 to identify the political tweets). Note that, these weights reduce the relevance of that terms that are recurrent in all the time-series. For each previously identified peak we create $tf-idf$ vectors for the tokenized news and tweets by employing the idf map, thus obtaining, for each day taken into account, two vector sets:

- \mathcal{N} the vectors' set of the news;
- \mathcal{T} the vectors' set of the tweets belonging to PTCorpus.

Subsequently we employed the cosine similarity between vectors to select the most correlated news w.r.t the peak taken into account as follows:

$$\arg \max_{n \in \mathcal{N}} \sum_{t \in \mathcal{T}} \frac{n \cdot t}{\|n\| \|t\|}$$

The results achieved are summarized in Table 6.

Finally, we have qualitatively compared the news identified with this approach with the trending topics of Twitter related to the day of each peak and we have noticed that the most of the news effectively correspond to one of the political daily trend. However, for few peaks, NData does not contain any news correlated with the majority of the tweets of that day. Looking at the Twitter trending topics, we can say that this happens whenever the political discussions on Twitter do not concern about facts reported in newspapers, such as discussions spontaneously grown in the Twitter community. A meaningful example concerns the trending topic #no2giugno: this movement asked for the suspension of the military parade of June the 2nd, seen as a waste of resources, to use these moneys to reconstruct the city of Emilia (Italian region) after the earthquakes of the 2012. This discussion generated two peaks (May the 30th and the 31st) that had not correlation with the traditional media news.

6 Conclusions and Future Works

In this work we analyse the well-known political attitude of political disaffection by using Twitter data through machine learning techniques; note that, to our knowledge, no studies have been proposed to analyze this political phenomenon employing Twitter. In order to validate the quality of the time-series extracted from Twitter data, we highlight the strong relation of these data with political disaffection as measured with public opinion surveys (measured through low intentions to vote and low political efficacy). Furthermore, we show that important political news of Italian newspapers are often correlated with the highest peaks of the produced time-series.

Note that, despite some works present many doubts about the possibility to perform an electoral prediction using Twitter data (Metaxas, Mustafaraj, and Gayo-Avello, 2011; Gayo-Avello, 2012), the different task of political disaffection, as suggested by the strong correlations between the time-series generated by employing the public opinion surveys and the Twitter data automatically extracted by our approach, could offer an interesting research topic for further

investigations. Moreover, it is important to notice that the Twitter timeliness to answer to a political events with respect to traditional public opinion surveys suggests that our method could be employed to perform a such of daily prediction of the the citizens political disaffection change.

In future works, we would like to extend our machinery in order to achieve better results. In particular we would like to integrate some interesting features from the technical point of view: the classification accuracy can be improved with a proper selection of the tweets to be labelled by experts in an active-learning fashion (i.e. Cesa-Bianchi, Gentile, and Orabona (2009)); furthermore, we would like to include the information about the graph topology in order to have a better understanding of the social component of these phenomena and the possibility to employ graph-based classifiers (i.e. Vitale et al. (2011), Herbster, Pontil, and Galeano (2008)).

References

- Adamic, L. A.; Baeza-Yates, R. A.; and Counts, S., eds. 2011. *Proceedings of the Fifth International Conference on Weblogs and Social Media, Barcelona, Catalonia, Spain, July 17-21, 2011*. The AAAI Press.
- Atkinson, R., and Flint, J. 2001. Accessing hidden and hard-to-reach populations: Snowball research strategies. *Social research update* 33(1):1–4.
- Birmingham, A., and Smeaton, A. 2011. On using twitter to monitor political sentiment and predict election results. In *Proceedings of the Workshop on Sentiment Analysis where AI meets Psychology (SAAIP)*. IJCNLP.
- Bollen, J., and Mao, H. 2011. Twitter mood as a stock market predictor. *IEEE Computer* 44(10):91–94.
- Bollen, J.; Mao, H.; and Pepe, A. 2011. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In Adamic et al. Adamic, Baeza-Yates, and Counts (2011).
- Breiman, L. 2001. *Random Forests*.
- Campbell, A.; Gurin, G.; and Miller, W. E. 1954. *The Voter Decides*. Evanston, IL: Row, Peterson and Company, 1st edition.
- Cesa-Bianchi, N.; Gentile, C.; and Orabona, F. 2009. Robust bounds for classification via selective sampling. *ICML*.
- Crammer, K.; Dekel, O.; Keshet, J.; Shalev-Shwartz, S.; and Singer, Y. 2006. Online passive-aggressive algorithms. *Journal of Machine Learning Research*.
- Di Palma, G. 1970. Apathy and participation. *Mass Politics in Western Societies* –.
- Gayo-Avello, D. 2012. "i wanted to predict elections. *CoRR* abs/1204.6441.
- Gentile, C. 2001. A new approximate maximal margin classification algorithm. *Journal of Machine Learning Research*.
- Gruhl, D.; Liben-Nowell, D.; Guha, R. V.; and Tomkins, A. 2004. Information diffusion through blogspace. *SIGKDD Explorations* 6(2):43–52.
- Herbster, M.; Pontil, M.; and Galeano, S. R. 2008. Fast prediction on a tree. *NIPS*.
- Krippendorff, K. H. 2003. *Content Analysis: An Introduction to Its Methodology*. Sage Publications, Inc, 2nd edition.
- Livne, A.; Simmons, M. P.; Adar, E.; and Adamic, L. A. 2011. The party is over here: Structure and content in the 2010 election. In Adamic et al. Adamic, Baeza-Yates, and Counts (2011).
- Lodhi, H.; Saunders, C.; Shawe-Taylor, J.; Cristianini, N.; and Watkins, C. J. C. H. 2002. Text classification using string kernels. *Journal of Machine Learning Research* 2:419–444.
- Manning, C. D.; Raghavan, P.; and Schütze, H. 2008. *Introduction to information retrieval*. Cambridge University Press.
- Metaxas, P. T.; Mustafaraj, E.; and Gayo-Avello, D. 2011. How (not) to predict elections. In *SocialCom/PASSAT*, 165–171. IEEE.
- O'Connor, B.; Balasubramanyan, R.; Routledge, B. R.; and Smith, N. A. 2010. From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. In *Proceedings of the International AAAI Conference on Weblogs and Social Media*.
- Orabona, F. 2009. *DOGMA: a MATLAB toolbox for Online Learning*. Software available at <http://dogma.sourceforge.net>.
- Popescu, A.-M., and Pennacchiotti, M. 2011. "dancing with the stars", nba games, politics: An exploration of twitter users' response to events. In Adamic et al. Adamic, Baeza-Yates, and Counts (2011).
- Rozza, A.; Lombardi, G.; Casiraghi, E.; and Campadelli, P. 2012. Novel fisher discriminant classifiers. *Pattern Recogn.* 45(10):3725–3737.
- Segatti, P. 2006. *Italy, forty years of political disaffection. A longitudinal exploration*. London: M. Torcal and J.R. Montero (eds). Political Disaffection in Contemporary Democracies: Social Capital, Institutions and Politics. Routledge. –.
- Shalev-Shwartz, S.; Singer, Y.; Srebro, N.; and Cotter, A. 2011. Pegasos: Primal estimated sub-gradient solver for svm. *Mathematical Programming, Series B*.
- Skoric, M.; Poor, N.; Achananuparp, P.; Lim, E.-P.; and Jiang, J. 2012. Tweets and votes: A study of the 2011 singapore general election. In *Proceedings of the 2012 45th Hawaii International Conference on System Sciences, HICSS '12*, 2583–2591. Washington, DC, USA: IEEE Computer Society.
- Tjong Kim Sang, E., and Bos, J. 2012. Predicting the 2011 dutch senate election results with twitter. In *Proceedings of the Workshop on Semantic Analysis in Social Media*, 53–60.
- Torcal, M., and Montero, J. 2006. *Political Disaffection in Contemporary Democracies: Social Capital, Institutions and Politics*. Routledge.

- Tumasjan, A.; Sprenger, T. O.; Sandner, P. G.; and Welp, I. M. 2010. Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*, 178–185.
- Vallina-Rodriguez, N.; Scellato, S.; Haddadi, H.; Forsell, C.; Crowcroft, J.; and Mascolo, C. 2012. Los twindignados: The rise of the indignados movement on twitter. *SocialCom/PASSAT*.
- Vitale, F.; Cesa-Bianchi, N.; Gentile, C.; and Zappella, G. 2011. See the tree through the lines: The shazoo algorithm. *NIPS*.
- Yelena Mejova, P. S. 2011. Exploring feature definition and selection for sentiment classifiers. In *In Fifth International AAAI Conference on Weblogs and Social Media (ICWSM)*.